

PROJECT NOTES

WIRELESS NETWORKS TO SUPPORT RAPID DECISION-MAKING

Assistant Professor Arnold Buss

Visiting Assistant Professor Paul Sanchez

Department of Operations Research

Introduction

The new IEEE 802.11 standard for wireless computing has sparked research into applying wireless commercial-off-the-shelf (COTS) technology to military problems. For the past year the Loosely Coupled Components (LCC) Working Group has been conducting research into how this emerging technology can best be leveraged to support military decision-making. Wireless computing offers the potential to deploy software systems which are much more dynamic than their predecessors. Our work addresses a number of questions that arise in the context of this emerging technology. The most significant issue is architectural. Wireless technology puts computing in the hands of a completely new set of users, users who are mobile and in a dynamic environment. We need to re-think the nature of the data that will be available and/or relevant, the algorithms, and the way in which results are presented. There are also issues of interoperability, bandwidth, range, and security.

The LCC Working Group has built a wireless network over the past year in Glasgow Hall to provide a "proof-of-concept" testbed. Drawing on this experience, we helped implement and support a wireless network for NPS' Center for Executive Education. These systems have been built on COTS devices and have utilized the 802.11 standard to promote interoperability between different manufacturers' devices as well as different operating systems. We believe that such heterogeneous networks are ultimately more useful for military applications than single-vendor turnkey solutions.

Architectures based on wireless computing are fundamentally much more dynamic than the more familiar wired technologies. However, many of the protocols that support computer networks are still rooted in static, wired networks. Most software systems, both commercial and DoD, have been designed with a single computer in mind. Only recently have the ideas of distributed computing, such as internet-based applications, begun to filter into software. Often the network elements of the design are added after the fact on pre-existing single platform applications. Even when designed with networking in mind, it is clear that most software developers had a static wired network model behind the design. Building applications that run on and exploit the capabilities of networks of mobile wireless is a research goal of the LCC group.

Research Goals

The LCC Working Group has worked for the past three years on new software architectures to support real-time decision-making. Inspired by *Joint Vision 2010*, the group has addressed problems relating to how software must be designed to take full advantage of both wide-area and local computer networks.

The process of exploiting these new technologies begins with the distinction between data and information. Most software planning systems, both COTS and DoD, have not made this distinction. Indeed, many applications appear to be directed towards how to display data to the user. It is clear that there will be an increasing abundance of data. It should also be clear that the implicit assumption that more data are automatically better for the decision-maker is false. There is a major risk that a decision-maker will be overwhelmed by too much data and will not be able to "see the forest for the trees."

The software architectures the LCC Working Group has been developing address the problem of data overabundance by applying Operations Research (OR) models to the data to produce information. This information is of much greater use to the decision-maker than the raw data.

The LCC architecture supports military applications in such basic tasks as the location of units, integration of intelligence information, and displaying data and information on maps. While technologies such as wireless computers and ubiquitous networks enable the fusion of existing and real-time data, more is required to support real-time decisions.

Some systems that have recently been developed using this architecture have enabled a decision-maker to quickly build graph and network models from real-time data for road, computer, or telecommunications networks and apply algorithms to these graphs. These algorithms enable the user to extract useful information, such as the shortest route between units' location and an objective [Bilyeu, 1998]; optimal arcs to interdict [Moriarty, 1997]; and which medical units should be matched with which tasks in a dynamic, rapidly changing environment [Bradford, 2000].

Last year the LCC Working Group provided support to a pair of SOLIC theses by LT Robert Moss, USN and LT Steven Tripp, USN [Moss, 1999; Tripp, 1999]. They explored some issues in wireless computing regarding range, bandwidth, and detectability of signal strength. Their work

--continued on page 19

PROJECT NOTES

WIRELESS NETWORKS, *continued from page 18*

represented our initial foray into wireless networking and provided a dramatic example of the power of wireless networks by carrying a wireless computer in an aircraft flying over the Monterey Bay. Using a directional antenna pointed from the aircraft to another antenna on the roof of Ingersoll Hall, they were able to maintain connectivity for up to 15 miles. From his office on campus, RADM Chaplin e-mailed to them a request for a picture of the Santa Cruz boardwalk while they were in the air. The desired image was received by e-mail less than 15 minutes later [Campus News, September 16, 1999].

Architecture

The LCC Working Group has been working on two types of wireless network implementations. One is a relatively static environment in which wireless clients arrive and leave but the

infrastructure is more or less fixed. This is the approach we have taken in Glasgow Hall. This quasi-static wireless network is appropriate when there is a need or desire to add the flexibility of wireless computing to existing wired networks. The result is a mixed network, consisting of both wired and wireless computers.

The second implementation is much more dynamic and comprises a wireless computer network that has an extremely small physical footprint and can be deployed very rapidly. It consists primarily of computers smaller than a notebook. Turning on these small units forms the network. It evolves and is reconfigured with each additional computer. As each device appears on the network, it “discovers” the existing devices and announces its presence to the network. We have dubbed this a “LAN in a Bag” because a complete network can be fit into a container the size of a conventional computer bag.

Currently the IEEE 802.11 standard requires the presence of an access point, a device that serves a role analogous to that of a hub in a wired Ethernet network. We have successfully used access points from two manufacturers: a Lucent WavePoint and an Apple AirPort. Currently the configuration in Glasgow Hall relies on both for coverage of the second floor, as shown in Figure 1.

The placement of the access points in Figure 2 provides nearly full coverage of the second floor, as well as partial coverage outside and on other floors. Note the striped block in Figure 1—this area is a storage closet with metal shelving and chain-link fence in the middle, and casts a significant radio “shadow.” This necessitates a greater density of access points in that portion of the building to guarantee adequate coverage. The use of Apple’s AirPort hubs is extremely cost-effective, since they are less than \$300 apiece. This compares quite favorably to Lucent’s equipment, at approximately \$1200 for a single WavePoint.

The access points provide gateways between the wireless clients and the wired LAN, as shown in Figure 2. The wired LAN provides access to the Internet, thus enabling the wireless clients to transparently utilize all of the Internet’s resources in addition to those of the wired LAN.

-- continued on page 20

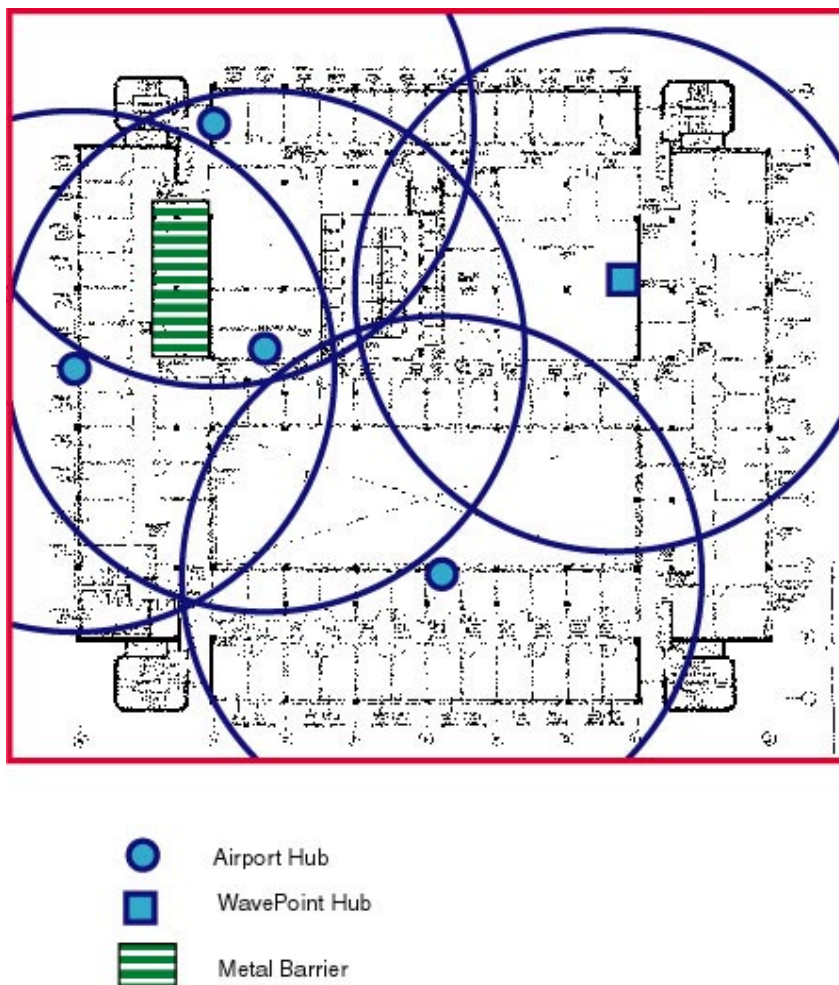


Figure 1. Access Point Coverage of Glasgow Hall 2nd Floor

PROJECT NOTES

WIRELESS NETWORKS, *continued from page 19*

Security

Security issues are a concern whether operating in a wired or wireless environment, but the broadcast nature of wireless computing brings them squarely to the forefront. Security issues can be broken into four broad categories: authentication, data protection, traffic patterns, and denial of service. The first two can be handled with a combination of encryption technology and/or Virtual Private Networks (VPNs). The latter two remain areas of vulnerability and are part of the ongoing research.

Current generation wireless equipment has hardware encryption available out of the box. It uses symmetric shared-key algorithms, and acts as the first layer of security. One cannot join the network without knowing the key. We advocate that all transmissions should also be software encrypted.

Authentication is often handled by a challenge/response system. A classic challenge/response system is the use of a password, but password schemes are most emphatically not sufficient for authentication. In an environment where information is broadcast through the airwaves using COTS technology, it must be assumed that message packets, and therefore passwords, will be intercepted. It follows from this that authentication must be based on a mechanism in which the response is never twice the same, but is nevertheless recognizable as correct. Public key encryption (PKE) technologies provide us with exactly such a mechanism. When a client attempts to make contact it must identify itself to the server. The server issues a challenge in the form of a randomly generated numeric value, which has been encrypted with the client's (known) public key. The client decrypts the number using his private key, then encrypts it with the server's (known) public key and returns it. A successful exchange establishes that both machines are who they claim to be, since to decrypt the number requires access to the private key that uniquely corresponds to the known public key. This eliminates both mimics (since the content of the challenge will change every time) and "man in the middle" attacks. It also has the potential to allow access rights to information to be restricted on an authorized user or authorized client basis, unlike the hardware encryption described above.

The random number that was exchanged for authentication purposes can be used as a key for symmetric key encryption of

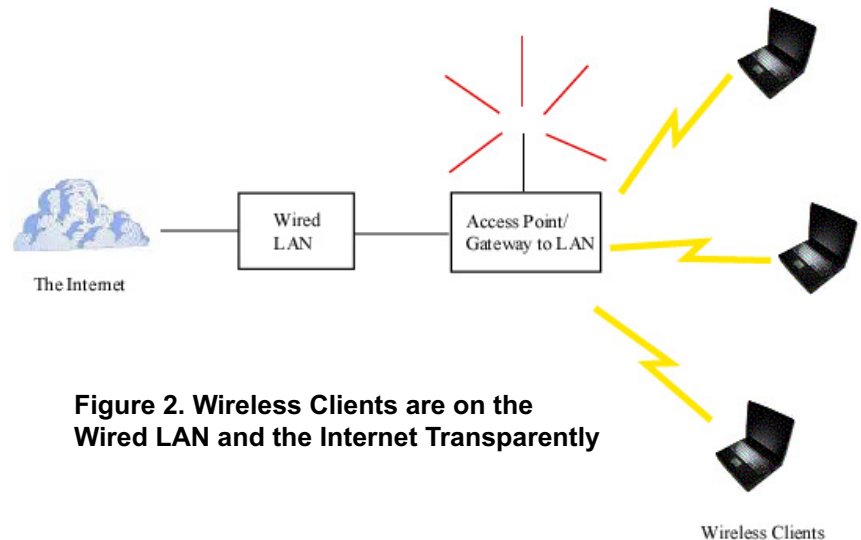


Figure 2. Wireless Clients are on the Wired LAN and the Internet Transparently

subsequent data. This is the mechanism used by VPN software to establish "secure tunnel" connections. Note that since the challenge is unique for each client session, the resulting encrypted channel is also unique (unlike the hardware encryption performed on the radio signals). This acts as a second layer of protection on the information content - even if the enemy has captured some of our equipment or succeeded in breaking the hardware encryption, they will not be able to view transmissions to other machines. Rather than assuming that we can keep the opposition out of our network infrastructure, we assume that they are present and make it impossible for them to view content regardless.

Another area of data vulnerability is from captured equipment. We are just beginning to investigate how technologies such as the Java-powered iButton can be used to prevent captured computers from being used by enemy forces. iButtons are miniature (about the size of a watch battery) computing devices which can be embedded in a ring, on a key chain, or mounted on a dog tag, and have the capability of doing public key encryption for authentication, as described above. With iButtons, smart-cards, or some equivalent technology, we anticipate the capability of authenticating the user to the computer itself. In other words, without the proper iButton and biometric data the computer would refuse to operate.

Oddly enough, the broadcast nature of wireless is somewhat helpful with regard to traffic pattern analyses. However, with sufficient effort it is likely that the opposition can detect

-- continued on page 21

PROJECT NOTES

WIRELESS NETWORKS, *continued from page 20*

acknowledgement responses and use these to identify patterns in transmission

Denial of service represents a potential vulnerability for wireless computing. The current generation of equipment uses spread spectrum and frequency hopping technologies but, as with any electromagnetic communications, is vulnerable to jamming.

Equipment

The Institute of Electrical and Electronics Engineers, Inc. (IEEE) created the 802.11 standard, which provides for vendor-independent interoperability of wireless devices. The original IEEE802.11 standard provided for transmission speeds of up to 2 megabits per second (Mb). An extended standard was adopted in late 1999 which provides for 11Mb communications. The channel can be either encrypted or unencrypted. Encryption is available at two security levels - 56 bit keys; or 128 bit keys, which provide greater security but are subject to export restrictions from the US.

Wireless networking in IEEE802.11 mode is topologically similar to wired networking using a hub. The access point must be connected to the internet via standard connections, such as ethernet or modem, if WAN (Wide Area Networking) connectivity is desired. Each client computer must have an IEEE802.11 compliant device, comparable to the NIC (Network Interface Card) used for ethernet connectivity. This is usually in the form of a PC card which fits into the standard PCMCIA slot available on virtually all laptops. Since desktop systems are usually viewed as non-mobile, wireless devices for them are much less common. However, ISA or PCI PCMCIA adapters are relatively inexpensive and enable desktop systems to join wireless networks. Regardless of whether they are laptop or desktop systems, client computers talk to the access point just as ethernet clients communicate with a hub, except that the communication is done by radio rather than by wire. Once properly configured, communications by wireless device are as transparent to the user as the wired connection currently on your desktop system.

In order to have "roaming" capabilities, all access points and clients must share the same IP address space, be configured to the same network name, and, if encryption is being used, share a common key. The range of the wireless connection is impacted by numerous factors, but we have found it to be dominated by distance between client and access point and the amount of metal (girders, wire mesh, metal bookcases, etc.) along the signal path. Our experience is that reception is

reasonable within a disk of radius ~150-300 ft centered on the access point. We have provided coverage for the entire second floor of Glasgow Hall using four access points.

In terms of specific vendor products, we are currently using a combination of Apple "Airport" and Lucent "Orinoco" technologies. Lucent currently offers three grades of cards, all operating at 11Mb — Gold cards can do 128b encryption; Silver cards do 56b encryption; and Bronze cards have no encryption capabilities. Any of these cards can be used for client connectivity or plugged into Lucent's hardware or a PC running Linux to act as an access point. At the time at which this article is being written, the cards cost ~\$200 each, while Lucent access points are ~\$1200 (without the required card). Apple makes the incredibly low-cost "AirPort base station," which sells for under \$300, contains a Lucent silver card and modem, and can provide both DHCP (Dynamic Host Configuration Protocol) and NAT (Network Address Translation) capabilities. There are two drawbacks to the airport. First, it can only be administered from a computer running MacOS. However, a quick calculation shows that an iMac + airport (~\$1300) is cheaper than an Orinoco access point + Wavelan card (~\$1400), and a price differential of ~\$1100 is realized with each additional access point. The second drawback is that the Apple airport contains a silver card, and is thus limited to 56-bit encryption. Since the key must be common to the entire network for roaming, this affects security for the entire system. However, this limitation is easily remedied with a screwdriver and a Lucent gold card. While this solution raises the cost of an Apple airport based system, it still is substantially cheaper than Lucent's solution.

It is worth pointing out that IEEE802.11 compliance provides us with vendor interoperability, and thus removes vendor-specific dependence. We have successfully built a network containing a combination of hardware and systems including Intel PCs, Lucent and Apple wireless equipment, and running various combinations of Windows, MacOS, Linux, FreeBSD, and OpenBSD. By focusing on open standards at the telecommunications level rather than at the application level, we have achieved seamless interoperability. This is significant for two reasons. First, the design allows for growth and evolution - individual components can be upgraded as improved versions become available as long as IEEE802.11 compatibility is maintained. Second, it increases the utility of the design for joint operation with allies or civilian organizations, such as the Red Cross or other rescue

-- continued on page 49

PROJECT NOTES

WIRELESS NETWORKS, *continued from page 21*

and aid organizations. Wireless communications seem like an obvious solution in situations such as the recent floods in Africa. Computing and networking can still be used even if wired telecommunications infrastructures have been damaged. If networking is based on IEEE802.11 protocols, the communications will work regardless of what computing platforms other agencies are using.

Campus Infrastructure

The LCC Working Group has been actively working to create an infrastructure for wireless networking which will eventually give users connectivity from any building (See Figure 3).

Campus Networking has committed a significant portion of a

Class C network for this purpose. This network, 131.120.14.*, provides the wired backbone for the wireless network. Wireless access points connected to this backbone provide connectivity to a common DHCP server, which allocates IP (Internet Protocol) addresses to the wireless clients. Since the DHCP server is common across the backbone, the leased IP address is valid for all access points along this subnet. The initial implementation was completed early this year on the second floor of Glasgow Hall. Students have found it particularly useful to be able to work from their study desks.

Campus network services can convert any specified ethernet port to be part of this effort. Once the ethernet port has been

switched to the common backbone, an access point of some sort needs to be connected. As mentioned earlier, we have found both Linux and Apple Airports to be cost-effective solutions. Finally, all of the client computers need to be configured with the shared network password, and to use DHCP. With a minimal investment of resources – an ethernet port dedicated to the wireless subnet and corresponding access point approximately every 200 feet – it would be possible to enable the entire campus for wireless networking.

Application

CPT Maximo Moore, USA, a student in the Operations Research Department and part of the LCC Working Group, has applied the LCC architecture running on a wireless network to Ranger Deployment Airfield Control Operations (DACO). The problem is to keep track of




Figure 3. Wired/Wireless Topology

--continued on page 50

PROJECT NOTES

WIRELESS NETWORKS, *continued from page 49*

personnel and equipment during a DACO. This is particularly important during the exfiltration part of the operation, where it is critical to identify soldiers or equipment that may have been loaded on different helicopters than they came in on.

CPT Moore developed a web-based wireless network solution using notebook computers. One laptop is designated a "server" – it must be running a database and web server application. The rest of the system is comprised of client laptops with barcode scanners and web browsers. Each soldier and piece of equipment will have a bar code that is scanned by a client laptop whenever he gets on or off any aircraft. The database containing information on where everyone (and everything) is located is updated wirelessly via a "servlet," a small application that is accessed using a standard web browser. The servlet can also ascertain the status of any individual or item of equipment. Servlets use standard HTML without any special requirements. Therefore, any device that is wireless-capable, authenticated to the network, and can run a browser can access the database using servlets.

This application illustrates the utility of the "LAN in a Bag" design. Deployment involves configuring all devices to a common network password, and then bringing the server and one or more clients within range of each other. The equipment is extremely portable, involves only COTS technologies, and the network that is the underpinning is completely transient and dynamic.

Other Uses

The LCC Working Group's wireless architecture was utilized recently by the Center for Executive Education (CEE) in a three-week course for high-ranking flag officers. The concept was that the participants could either bring their notebook computers and be issued a wireless card or use one of the Center's notebook computers for the course's computing requirements. The wireless computers were critical to the conduct of the module. For example, small breakout groups had complete internet access no matter where they met within the Center. All participants were able to communicate with each other and share documents on the computers.

The network also illustrated the wired/wireless concept as well as mixed network ideas. In order to obtain internet access, both a WavePoint and an AirPort hub were used. Connectivity was seamless, and the users were unaware of which access point was being used at any time. Furthermore,

an old 486 computer was rescued from excess and configured with the Linux operating system to issue IP addresses using DHCP. When a wireless notebook was turned on and the card inserted, it issued a request for an IP address that went onto the wired LAN from either the AirPort or the WavePoint hub. This request was captured by the 486 Linux machine, which issued the address. From that point on the client had full IP networking, including internet access.

Ongoing Work

The LCC Working Group is continuing to work on various ways that new and emerging computing technologies can best be exploited to improve military planning and decision-making. The focus is on how these technologies can be best exploited to improve military planning and decision-making using tools from Operations Research and the Decision Sciences.

The research has been sponsored by grants from the Air Force Office of Scientific Research (AFOSR), NPS' Institute for Joint Warfare Analysis (IJWA), Office of Naval Research (ONR), and USSOCOM.

References

- [1] Bilyeu, Allan L., "Concept for a Special Operations Planning and Analysis System," MS in Operations Research, June 1998.
- [2] Bradford Robert, "Solving Dynamic Battlespace Movement Problems Using Dynamic Distributed Computer Networks," MS in Operations Research, June 2000.
- [3] Curley, PH2 Larry, "Wireless LAN Soars Through the Clouds," *Campus News*, September 16, 1999.
- [4] Moriarty, Sean T., "A Technical Demonstration of a Map Based Logistics Planning Tool," MS in Operations Research, September 1997.
- [5] Moss, Robert B., "SOF Tactical Intranet : Low Probability of Detection, Low Probability of Exploitation Communications for Special Operations Forces, Using A Commercial Off The Shelf Wireless Local Area Network," M.S. in Systems Technology and M.S. in Defense Analysis, September 1999.
- [6] Tripp, Stephen J., "An Airborne High Data Rate and Low Cost Digital Communications Network Using Commercial Off The Shelf Wireless Local Area Network Components," M.S. in Systems Engineering, September 1999.